Author Names & Affiliations

- James Stone - Princeton University
- Anatoly Spitkovsky - Princeton University
- Adam Burrows - Princeton University
- Matt Kunz - Princeton University
- Eve Ostriker - Princeton University

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Astrophysics, plasma astrophysics, MHD

Title of Submission

Cyberinfrastructure needs for plasma astrophysics in the coming decade

Abstract (maximum ~200 words).

Plasma astrophysics seeks to solve some of the most fundamental questions in the universe, such as how planets, stars, and galaxies form, how non-thermal energetic particles (cosmic rays) are produced, and how magnetic fields are created and amplified in astrophysical systems. The complex physics and nonlinear dynamics of astrophysical plasmas mean numerical methods are the most important tool for investigating these problems. Support for a broad range of computational facilities is needed, from local mid-scale developmental systems at universities and research institutes, to reliable and substantive resources at national facilities such as TACC, to access to the largest multi-petaflop (and indeed exaflop) systems for the largest users. Significant support for people is also required, including students and postdocs to develop algorithms, and software engineers to help implement and migrate codes to emerging architectures.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

Most of the visible matter in the universe is a plasma, and understanding the structure and dynamics of astrophysical plasmas is one of the major challenges of the field for the next decade. Basic problems such as understanding how galaxies form from density perturbations in the early universe, how stars form from the turbulent, magnetized interstellar medium in galaxies, and how planets form in dense cold accretion disks around young stars are all challenges in plasma astrophysics. Currently all of these problems are being addressed by multiple research groups, but further progress requires the addition of more physics to make calculations more realistic, and increased dynamic range in scales and time to improve the accuracy and fidelity of the results.

For dense collisional plasmas, the dynamics is well described by the equations of magnetohydrodynamics (MHD), however for partially ionized plasmas non-ideal effects such as resistivity, ambipolar diffusion, and the Hall effect must be included. Additional physics such as self-gravity, optically thin cooling, radiation transport, neutrino transport, non-ideal gas equation of state, and chemistry must be included in many cases. The multi-physics nature of many problems makes them very challenging. Again, multiple research groups currently study the non-ideal MHD of planet formation, the radiation-dominated MHD of accretion of matter into black holes, and the complex physics of core-collapse supernovae. However, across the board, progress is limited by the software used for these problems, and by insufficient computational resources.

For diffuse plasmas, a fully kinetic approach using particle-in-cell (PIC) or hybrid-PIC methods must be adopted. Important problems in these regimes are understanding the acceleration of high-energy particles in collisionless shocks, cosmic ray diffusion and feedback in galaxies, and the dynamics of diffuse accretion flows such as at the center of our own galaxy.

An emerging and promising area for the future are multiscale methods, where the dynamics is modeled using MHD on large scales, and PIC or hybrid PIC on small scales to capture kinetic effects properly. Such approaches need adaptive meshes with very high dynamic range, and special methods to smoothly transition between regimes. They are beginning to be applied to space physics problems, and show much promise for the future. Development and deployment of these methods will be very labor intensive.

Finally, there is increasing overlap between space physics, laboratory plasma physics, and astrophysics. The same methods and problems appropriate for studying the heliosphere apply to many problems in astrophysics. Similarly, the techniques used to study astrophysical plasmas can be applied to model laboratory reconnection and dynamo experiments, and collisionless shock experiments performed at high-energy density laser facilities. There is the potential for tremendous synergy in these areas.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

Infrastructure required to support this research spans the entire computing ecosystem.

• Local, mid-scale clusters dedicated to code development and testing are essential. Programs that support universities and research institutes to acquire such mid-scale systems (e.g. through MRI program) is important

• Access to a wide spectrum of national HPC resources is crucial. For the majority of research groups, access to tens of millions of core-hours at national facilities such as TACC is the foundation of their research program. At the same time, for a small but important set of research projects, access to hundreds of millions of core-hours for well-resolved runs including complex physics at the largest scale such as Blue Waters is essential.

In addition to hardware, support for developing the algorithms and codes needed to tackle research problems in astrophysical plasma physics is essential. Such support always seems to lag the investment in hardware, but obviously effective use of emerging hardware is impossible without code development. Moreover, adding more physics to existing codes to enable research at the edge of knowledge is hindered without adequate support for people. Fellowships for students and postdocs for algorithm development, and for software engineers who can help profile, tune, and enhance performance codes on modern architectures is needed.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Over the past decade the DOE has provided the best support for high-end computing for science and engineering through programs like INCITE. Ideally, the NSF should be partnering with the DOE to support computational science in a more effective manner. The DOE has a very ambitious plan to achieve exascale by 2021, and it seems the NSF is playing very little role in partnering with other agencies in providing exascale to the research community. At the very least, it should provide domain-specific support for development of codes (e.g. by supporting people to develop such codes) for exascale.

In astrophysics, the emergence of community codes that are widely used by thousands of researchers has accelerated progress. The NSF should consider ways it can more directly support the development and maintenance of community codes.

Consent Statement